

Although FIGS. 3A and 3B illustrate exemplary embodiments of deterministic resource management having an existing cycle time T and read-ahead size N_i that are modified as necessary to support new or returning viewers, it will be understood that a variety of other embodiments using different methodology for resource management are also possible using the disclosed methods and systems. Such embodiments may be based at least in part on the methodology described in FIGS. 3A and or 3B, and/or described elsewhere herein. For example, resource management may be accomplished by determining or modifying only cycle time T or only read-ahead size N_i . Furthermore, it is not necessary that admission control decisions be made in conjunction with such determinations or modifications. Alternatively, it is possible that admission control decisions may be made without any determination or modification of information management system I/O operational parameters. In yet other embodiments, initial values of information management system I/O operational parameters such as cycle time T and/or read-ahead size N_i may be determined in addition to, or as an alternative to, later modification of such information management system I/O operational parameters.

In yet other embodiments, cycle time may be modified or limited based on a number of factors. For example cycle time may be limited or capped by limiting read-ahead buffer size, for example, using Resource Model Equations (17B), (18B) or (19B). Cycle time may also be limited or capped by placing a set limit on the maximal buffer size (e.g., by placing a 2MB limit on the maximal buffer size in a case where system throughput does not increase, or does not increase significantly, with any increase in the buffer size beyond 2MB).

Monitoring of System I/O performance characteristics

In the practice of the disclosed methods and systems, monitored system I/O performance characteristics may be at least partially considered or employed to effect information management system I/O resource management actions, for example, by using a resource model embodiment such as described elsewhere herein.

FIG. 4A illustrates one embodiment of a storage system 400 having a storage management processing engine 410 (e.g., storage processor) that includes resource manager 420, logical volume manager 430 and monitoring agent 440. In this embodiment, resource manager 420 is responsible for cache memory management, I/O admission control and resource monitoring and may therefore include cache memory manager 422, I/O admission controller 424 and storage system workload monitor 426. Storage management processing engine 410 may include any hardware configuration e.g., configuration of one or more processors or processing modules, that is capable of performing monitoring, resource modeling, resource management, and/or storage management duties described herein.

10

As shown in FIG. 4A, storage devices 450 may be coupled to storage management processing engine 410 by way of, for example, fiber channel loop or other suitable method. It will be understood that each storage device 450 may be a single storage device (e.g., single disk drive) or a group of storage devices (e.g., partitioned group of disk drives), and that combinations of single storage devices and storage device groups may be coupled to storage management processing engine 410. In one embodiment, storage management processing engine 410 may include one or more Motorola POWER PC-based processor modules. In the embodiment of FIG. 4A, it will be understood that storage devices 450 (e.g., disk drives) may be controlled at the disk level by storage management processing engine 410, and/or may be optionally partitioned into multiple sub-device layers (e.g., sub-disks) that are controlled by single storage processing engine 410.

Examples of these and other suitable storage management processing engine configurations, as well as examples of information management system environments in which storage management processing engines may be implemented in the practice of the disclosed methods and systems include, but are not limited to, those described in co-pending United States patent application serial number 09/797,413 filed on March 1, 2001 which is entitled NETWORK CONNECTED COMPUTING SYSTEM; in co-pending United States patent application serial number 09/797,200 filed on March 1, 2001 which is entitled SYSTEMS AND METHODS FOR THE DETERMINISTIC MANAGEMENT OF INFORMATION; and in co-pending United States patent application serial number 09/879,810 filed on June 12, 2001 which

is entitled SYSTEMS AND METHODS FOR PROVIDING DIFFERENTIATED SERVICE IN INFORMATION MANAGEMENT ENVIRONMENTS; each of the foregoing applications being incorporated herein by reference. Other examples include, but are not limited to, in an external RAID controller configuration, *etc.*

5

In the embodiment of FIG. 4A, workload on each storage device 450 may be logically monitored in a collaborative manner between resource manager 420 and logical volume manager 430 using, for example, workload monitor 426 and monitoring agent 440. In this regard, monitoring agent 440 may be employed to track the outstanding I/O requests in each storage 10 device, and workload monitor 426 of resource manager 420 may be employed to track the number of viewers and the aggregated playback rates for each logical volume. In this embodiment, resource manager 420 has the knowledge of the number of plex (*i.e.*, number of portions of a logical volume per storage device) in each logical volume, and therefore the total number of viewers and the aggregated playback rates on each logical volume may be averaged across plex to obtain the estimation of the total number of viewers and the aggregated playback 15 rates on each plex. Thus, workloads may be logically monitored or tracked above the logical volume manager level in terms of number of streams and aggregated playback rates. Advantageously, the disclosed methods and system may be implemented to monitor resource utilization and workload distribution at the logical volume level rather than the physical disk 20 level. For example a combination of logical volume workload, the logical volume level topology, and the number of outstanding I/O commands on each logical subdisk may be used to estimate workload characteristics at the physical disk level.

If, in a particular logical volume, there is only one storage device 450 (*e.g.*, disk drive) 25 per plex, then workload monitor 426 of resource manager 420 may have knowledge of workloads on each individual disk drive without assistance from monitoring agent 440. For example, this is the case as described herein in Examples 7, 9 and 10, where the storage organization may be one disk drive worth of content with several mirrors. However, if there is more than one storage device 450 (*e.g.*, disk drive) per plex, then the workload in the plex level 30 may be refined to obtain the workload view at the disk drive level. For example, this is the case